

REMARKS

In the non-final Office Action, the Examiner rejected claims 1-7 and 12-18 under 35 U.S.C. § 102(b) as anticipated by Bolosky et al. (U.S. Patent No. 5,991,804); rejected claims 8 and 9 under 35 U.S.C. § 103(a) as unpatentable over Bolosky et al. in view of Frey et al. (U.S. Patent No. 6,725,392); and rejected claims 10 and 11 under 35 U.S.C. § 103(a) as unpatentable over Bolosky et al. in view of Jacobs et al. (U.S. Patent Application Publication No. 2003/0023898).

By this Amendment, Applicants amend claim 4 to improve form. Applicants respectfully traverse the Examiner's rejections under 35 U.S.C. §§ 102 and 103. Claims 1-18 remain pending.

REJECTION UNDER 35 U.S.C. § 102 BASED ON BOLOSKY ET AL.

In paragraph 5 of the Office Action, the Examiner rejected claims 1-7 and 12-18 under 35 U.S.C. § 102(b) as allegedly anticipated by Bolosky et al. Applicants respectfully traverse the rejection.

A proper rejection under 35 U.S.C. § 102 requires that a single reference teach every aspect of the claimed invention either expressly or impliedly. Any feature not directly taught must be inherently present. In other words, the identical invention must be shown in as complete detail as contained in the claim. See M.P.E.P. § 2131. Bolosky et al. does not disclose or suggest the combination of features recited in claims 1-7 and 12-18.

For example, independent claim 1 is directed to a file system that comprises a plurality of servers configured to store data and a master connected to the servers. The master is configured to communicate with the servers upon startup of the master to authoritatively identify the data

stored by the servers, and record location information that identifies ones of the servers that store the data.

Bolosky et al. does not disclose or suggest the combination of features recited in claim 1. For example, Bolosky et al. does not disclose or suggest a master that is configured to communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers. The Examiner alleged that Bolosky et al. discloses this feature and cited column 2, lines 55-57, and column 16, lines 34-48, of Bolosky et al. for support (Office Action, page 3). Applicants respectfully disagree.

Before addressing the sections identified by the Examiner, it may be useful to provide some context to them. Bolosky et al. discloses a restriping method in which data blocks are moved around within the file server following a capacity change in which one or more storage disks and/or data servers are added or removed (col. 6, lines 35-41). Bolosky et al. discloses that the restriping method includes three phases: a first phase in which controller 22 builds an ordered work list for each disk 28 supported by the data servers 24 (where the work list describes the data block moves that the data servers 24 are to perform) and distributes the work lists to respective data servers 24; and second and third phases performed by data servers 24 in parallel to execute the data block moves in the work lists (col. 6, lines 45-60). Bolosky et al. discloses that data servers 24 use a logging scheme that reflects their progress in completing their work lists (col. 16, line 34 - col. 17, line 67).

At column 2, lines 55-57, Bolosky et al. discloses "In the event of system failure, the system restarts at the log pointer and proceeds from there." In this section, Bolosky et al. discloses that in the event of a system failure, data servers 24 (which the Examiner equated to the

plurality of servers in claim 1) restart execution of their work lists based on their logs (see also, col. 16, line 34 - col. 17, line 67). Nowhere in this section, or elsewhere does Bolosky et al. disclose or remotely suggest that controller 22 (which the Examiner equated to the master in claim 1) is configured to communicate with the servers upon startup of controller 22 to authoritatively identify the data stored by the servers, as required by claim 1.

At column 16, lines 34-48, Bolosky et al. discloses:

The second category of failures--whole system failures--is handled using a logging scheme. In the event of a system crash, such as via power failure, the restriping process is restarted at the last log entry. The logging scheme is rooted in each storage disk tracking which of its contents have been written to other disks, and which of the contents for which it holds the mirrored redundancy piece 0 have been written. Because the restriping process is performed in the order specified in the work list, keeping track of progress involves tracking a single work list index for the primary data block and another index for the first mirrored redundancy piece. These indices are stored in the header block of the disk in pre-reserved locations. The work list itself is also stored, either in the header block of the storage disk, or on the boot disk.

In this section, Bolosky et al. discloses that in the event of a whole system failure, data servers 24 restart the restriping process at the last log entry. Nowhere in this section, or elsewhere does Bolosky et al. disclose or remotely suggest that controller 22 (which the Examiner equated to the master in claim 1) is configured to communicate with the servers upon startup of controller 22 to authoritatively identify the data stored by the servers, as required by claim 1.

The Examiner also alleged "during whole system failure, it can be concluded that the controller fails; when the system restarts, the controller uses the log to determine the data stored by the servers in order to continue the transfer process" (Office Action, page 3). Applicants respectfully submit that the Examiner has misconstrued the disclosure of Bolosky et al. As explained above, Bolosky et al. discloses that data servers 24, and not controller 22, use the logs to restart the restriping process at the last log entry (col. 16,

lines 34-48). Therefore, the Examiner's allegation lacks merit.

Even assuming, for the sake of argument, that the Examiner's allegation is accurate (a point that Applicants do not concede), Applicants submit that the Examiner is not addressing the features of claim 1. Claim 1 does not recite that the master uses a log to determine data stored by the servers. Instead, claim 1 recites that the master is configured to communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers. Therefore, the Examiner's allegation falls short of establishing a proper case of anticipation.

Because Bolosky et al. does not disclose or suggest a master that is configured to communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers, Bolosky et al. cannot disclose or suggest that the master is further configured to record location information that identifies ones of the servers that store the data, as further recited in claim 1.

For at least these reasons, Applicants submit that claim 1 is not anticipated by Bolosky et al. Claims 2-7 and 12 depend from claim 1 and are, therefore, not anticipated by Bolosky et al. for at least the reasons given with regard to claim 1.

Independent claim 13 is directed to a master in a file system that includes a master connected to a plurality of servers. The master comprises means for performing a startup operation; means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers; and means for storing location information that identifies ones of the servers that store the data.

Bolosky et al. does not disclose or suggest the combination of features recited in claim

13. For example, Bolosky et al. does not disclose or suggest a master that comprises means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers. The Examiner alleged that Bolosky et al. discloses this feature and cited column 3, lines 43-45, and item 26 in Fig. 1, of Bolosky et al. for support (Office Action, page 6). Applicants respectfully disagree.

At column 3, lines 40-46, Bolosky et al. discloses:

The file server 20 might alternatively operate as a content provider that distributes data files over a network (e.g., Internet, LAN, etc.) to multiple client computers.

The continuous media file server 20 has a central controller 22 connected to multiple data servers 24(1), 24(2), 24(3), . . . , 24(K) via a low bandwidth control network 26.

In this section, Bolosky et al. discloses that controller 22 connects to data servers 24 via a low bandwidth control network 26. Nowhere in this section, or elsewhere, does Bolosky et al. disclose or remotely suggest a master that comprises means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers, as required by claim 13.

Bolosky et al. describes item 26 in Fig. 1 as a low bandwidth control network that connects controller 22 to data servers 24 (col. 3, lines 44-46). Nowhere in regard to item 26, or elsewhere, does Bolosky et al. disclose or remotely suggest a master that comprises means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers, as required by claim 13.

If this rejection is maintained, Applicants respectfully request that the Examiner explain how the above sections of Bolosky et al. can reasonably be construed to disclose or suggest a master that comprises means for communicating with the servers during or after the startup

operation to authoritatively identify the data stored by the servers, as required by claim 13.

For at least these reasons and the reasons given above with regard to claim 1, Applicants respectfully submit that claim 13 is not anticipated by Bolosky et al.

Independent claim 14 is directed to a method for maintaining data in a file system that includes a master connected to a plurality of servers. The method, performed by the master, comprises communicating with the servers upon startup of the master to authoritatively determine data stored by the servers; and generating location information based on the data determined to be stored by the servers.

Bolosky et al. does not disclose or suggest the combination of features recited in claim 1. For example, Bolosky et al. does not disclose or suggest communicating, by the master, with the servers upon startup of the master to authoritatively determine data stored by the servers.

The Examiner alleged that Bolosky et al. discloses this feature and cited column 2, lines 55-57, and column 16, lines 34-48, of Bolosky et al. for support (Office Action, page 6). For at least reasons similar to reasons given with regard to claim 1, Applicants respectfully submit that neither of the sections identified by the Examiner discloses or remotely suggests communicating, by the master, with the servers upon startup of the master to authoritatively determine data stored by the servers, as required by claim 14.

For at least these reasons, Applicants respectfully submit that claim 14 is not anticipated by Bolosky et al.

Independent claim 15 is directed to a file system that comprises a plurality of servers configured to store files as chunks and a master connected to the servers. The master is configured to authoritatively determine location information by communicating with the servers,

the location information being based on which of the servers store ones of the chunks, and update the location information by periodically communicating with the servers to obtain changes to the location information.

Bolosky et al. does not disclose or suggest the combination of features recited in claim 15. For example, Bolosky et al. does not disclose or suggest a master that is configured to authoritatively determine location information by communicating with the servers, where the location information is based on which of the servers store ones of the chunks. The Examiner alleged that Bolosky et al. discloses this feature and cited column 3, lines 43-45, and column 5, lines 1-3, of Bolosky et al. for support (Office Action, page 7). Applicants respectfully disagree.

Column 3, lines 43-45, has been reproduce above. In this section, Bolosky et al. discloses that controller 22 connects to data servers 24 via a low bandwidth control network 26. Nowhere in this section, or elsewhere, does Bolosky et al. disclose or suggest a master that is configured to authoritatively determine location information by communicating with the servers, where the location information is based on which of the servers store ones of the chunks, as required by claim 15.

At column 5, lines 1-3, Bolosky et al. discloses ". . . a map can be maintained at the controller 22 or in the data servers 24 to track the physical locations of the data blocks." This section of Bolosky et al. discloses that controller 22 can maintain a map that tracks the physical locations of the data blocks. Nowhere in this section, or elsewhere, does Bolosky et al. disclose or suggest that controller 22 (which the Examiner alleged was equivalent to the master in claim 15) authoritatively determines location information by communicating with the servers, where the location information is based on which of the servers store ones of the chunks, as required by

claim 15.

Bolosky et al. also does not disclose or suggest a master that is configured to update the location information by periodically communicating with the servers to obtain changes to the location information, as further recited in claim 15. The Examiner alleged that Bolosky et al. discloses this feature and cited column 9, line 64 - column 10, line 14, of Bolosky et al. for support (Office Action, page 8). Applicants respectfully disagree.

At column 9, line 64 - column 10, line 14, Bolosky et al. discloses:

In phase three of the restriping process, the data servers operate in parallel to move the data blocks. The assignments of data blocks to specific destination disk locations helps establish a particular order for executing the block moves so that data blocks presently residing in locations are safely moved before being overwritten by data blocks that are destined to occupy the same locations. The data block assignments can be realized as block movement summaries that become entries on the work list. There is one block movement summary for each data block. A block movement summary specifies a source location on a source disk where a particular data block presently resides and a destination location on a destination disk where the particular data block is to be moved. The block movement summary is represented as follows:

<Destination Disk, Destination Location>.rarw.<Source Disk, Source Location >

In this section, Bolosky et al. discloses steps involved in phase three of the restriping process, where the data servers operate in parallel to move the data blocks according to the work lists provided by the controller. Nowhere in this section, or elsewhere, does Bolosky et al. disclose or suggest a master that is configured to update the location information by periodically communicating with the servers to obtain changes to the location information, as required by claim 15.

For at least these reasons, Applicants submit that claim 15 is not anticipated by Bolosky et al.

Independent claims 16-18 recite features similar to, but possibly different in scope from,

features recited in claim 15. Claims 16-18 are, therefore, not anticipated by Bolosky et al. for at least reasons similar to reasons given with regard to claim 15.

REJECTION UNDER 35 U.S.C. § 103 BASED ON BOLOSKY ET AL. AND FREY ET AL.

In paragraph 7 of the Office Action, the Examiner rejected claims 8 and 9 under 35 U.S.C. § 103(a) as allegedly unpatentable over Bolosky et al. in view of Frey et al. Applicants respectfully traverse the rejection.

Claims 8 and 9 ultimately depend from claim 1. Without acquiescing in the Examiner's rejection with regard to claims 8 and 9, Applicants respectfully submit that the disclosure of Frey et al. does not cure the deficiencies in the disclosure of Bolosky et al. identified above with regard to claim 1. For example, Frey et al. is directed to a controller fault recovery system for recovering from faults that cause unscheduled stops for a distributed file system operating on an array storage system having multiple controllers (col. 4, lines 36-39). Frey et al. does not disclose or suggest a master that is configured to communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers, as required by claim 1. Therefore, claims 8 and 9 are patentable over Bolosky et al. and Frey et al. for at least the reasons given with regard to claim 1.

For at least these reasons, Applicants submit that claims 8 and 9 are patentable over Bolosky et al. and Frey et al., whether taken alone or in any reasonable combination.

REJECTION UNDER 35 U.S.C. § 103 BASED ON BOLOSKY ET AL. AND JACOBS ET AL.

In paragraph 8 of the Office Action, the Examiner rejected claims 10 and 11 under 35 U.S.C. § 103(a) as allegedly unpatentable over Bolosky et al. in view of Jacobs et al. Applicants respectfully traverse the rejection.

Claims 10 and 11 ultimately depend from claim 1. Without acquiescing in the Examiner's rejection with regard to claims 10 and 11, Applicants respectfully submit that the disclosure of Jacobs et al. does not cure the deficiencies in the disclosure of Bolosky et al. identified above with regard to claim 1. For example, Jacobs et al. is directed to a method for replicating data from a master server to at least one slave or managed server (para. 0009). Jacobs et al. does not disclose or suggest a master that is configured to communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers, as required by claim 1. Therefore, claims 10 and 11 are patentable over Bolosky et al. and Jacobs et al. for at least the reasons given with regard to claim 1.

For at least these reasons, Applicants submit that claims 10 and 11 are patentable over Bolosky et al. and Jacobs et al., whether taken alone or in any reasonable combination.

CONCLUSION

In view of the foregoing amendments and remarks, Applicants respectfully request the Examiner's reconsideration of the application and the timely allowance of pending claims 1-18.

If the Examiner does not believe that all pending claims are now in condition for allowance, the Examiner is urged to contact the undersigned to expedite prosecution of this application.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, L.L.P.

By: /Paul A. Harrity/
Paul A. Harrity
Reg. No. 39,574

Date: April 24, 2006

11350 Random Hills Road
Suite 600
Fairfax, Virginia 22030
(571) 432-0800